(Preprint) AAS 19-456

SIMULTANEOUS OPTIMIZATION OF SPACECRAFT AND TRAJECTORY DESIGN FOR INTERPLANETARY MISSIONS UTILIZING SOLAR ELECTRIC PROPULSION

Austin K. Nicholas*, Ryan C. Woolley†,

Alan Didion‡, Frank Laipert§, Zubin Olikara, Ryan Webb††, Rob Lock‡‡**

A major challenge in formulating interplanetary mission concepts utilizing electric propulsion is the large number of trajectory variables that must be considered (thrust profile, flyby options, launch vehicle delivery), all of which are affected by spacecraft design variables (power, mass, thruster, payload, staging). This is significantly more complex than traditional ballistic/chemical mission design and early concepts are often suboptimal as a result, potentially missing valuable options. This paper presents a novel tool (MORT) for simultaneously optimizing the spacecraft design alongside the trajectory given mission constraints and objectives, including example results relevant to the exploration of Mars.

INTRODUCTION

In the last decade, electric propulsion (EP) has seen widespread adoption in both the commercial market and in exploration missions. Electric propulsion is changing the paradigm of deep space missions and opening up opportunities that would be impossible or impractical with only chemical propulsion. This is well-illustrated by the recent success of NASA's Dawn mission, which was able to orbit two target bodies (Vesta, Ceres) in the same mission. Upcoming missions, including BepiColombo, Psyche and numerous other proposed concepts also include the use of solar electric propulsion (SEP). There is also interest in the human exploration community in using SEP for cargo missions to Mars[1] and at a potential lunar gateway. There is a rich landscape for EP concepts in the coming decades.

---

* Systems Engineer, Project Systems Engineering and Formulation
† Mission Design Engineer, Inner-Planets Mission Analysis
‡ Systems Engineer, Project Systems Engineering and Formulation
§ Mission Design Engineer, Flight Path Control
** Mission Design Engineer, Inner-Planets Mission Analysis
†† Systems Engineer, Project Systems Engineering and Formulation
‡‡ Manager, Mars Program Formulation Office
All Authors: NASA Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

Traditionally, when formulating ballistic deep-space missions, the primary consideration has been of the so-called "porkchop plots" - solutions to the Lambert problem. For more challenging targets, gravitational assists may be added and chained together to improve delivered mass, perhaps slightly modified using deep-space maneuvers. Overall, these types of trajectories have little to do with the spacecraft and can largely be computed in advance, independent of the mission and catalogued[2] for later retrieval.

For EP missions, on the other hand, the spacecraft propulsion performance is integral to the feasibility of the trajectory. Specifically, the spacecraft mass, thrust level, specific impulse (Isp), and power system all play critical roles in the trajectory performance. Further, for SEP missions, the performance also varies as a function of range (due to solar flux) and time (due to array degradation). Because of the large orbit changes possible with the high Isp of EP systems, the launch targets are also subject to optimization and less obvious from Lambert solutions. Flybys may still be helpful, and are possible under conditions not possible ballistically. These factors make EP trajectories largely spacecraft/mission-specific.

To complicate things even more, the dry mass of an SEP spacecraft is strongly dependent on the power and electrical propulsion components. As a result, trades involving trajectory performance are usually not straightforward. For instance, it is often not clear whether the performance gain from increasing array size offsets the increased weight for a given set of constraints. This expansive trade space for EP mission is difficult to explore and optimize due to the complex, non-linear coupling of these numerous factors[3].

The desire to explore the utility of a SEP orbiter to Mars[4], in view of the previously noted difficulties, led to the development of the **M**ars **OR**biter **T**ool, or **MORT**. The primary goal of MORT is to quantitatively explore the various in-space transportation options available to Mars missions and assess compatibility with a variety of possible mission objectives. The primary development of MORT occurred in 2014 to 2018 alongside the Mars formulation activities being undertaken in the Mars Program Formulation Office at NASA's Jet Propulsion Laboratory (JPL). MORT has been used to explore a diversity of concepts, including remote science orbiters, communications assets, sample return missions (from the moons and surface) as well as crewed and crew-cargo delivery. It has also been used in conjunction with a variety of mission modes, including direct launch, rideshare, and re-use (e.g. refueling at lunar depot). MORT also spans a variety of mission sizes, including SmallSats of order 100 kg, planetary robotic missions on the order of a few 1000's of kg, and crewed missions up to 100 metric tons.

MORT OVERVIEW

At its most basic essence, MORT is intended to ensure that the spacecraft and its associated trajectory are simultaneously feasible. A primary concern for all space missions is of course to ensure that the launch vehicle mass capability is not exceeded and that the tanks are sized properly to carry the propellant needed to perform all maneuvers. For SEP missions, it is critical that the onboard propulsion capability is compatible with the desired trajectory from an acceleration perspective for all legs. For Chemical Propulsion (CP) missions, it is important to consider gravity losses around impulsive maneuvers and also the impact of mass and drag area on aerobraking (if used).

The approach taken in developing MORT is to separate the spacecraft sizing problem from the trajectory problem and solve them iteratively to convergence. The rationale for this particular separation is largely to leverage the existing tools already available for trajectory optimizations. Figure 1 shows a simplified block diagram of the MORT modules as configured for an example round-trip Mars mission. Within the main spacecraft module, there are sub-modules for the relevant subsystems. Within the main trajectory module, there are sub-modules for each leg and/or significant maneuver. The information passed between modules depends on the specific problem at hand and the desired fidelity. In some cases it can be as simple as shown, in other cases there may be hundreds of parameters passed between the various modules.

At the heart of MORT is the simultaneous optimization of both the spacecraft module and the trajectory module. During intermediate iterations of the tool, the spacecraft module provides mass, power, and propulsion inputs at each trajectory leg into the trajectory module. The trajectory module returns propellant and times of flight for each leg, which then adjust the spacecraft sizing. This process can be completed either forwards or backwards in time, depending on mission constraints, until a feasible, optimized mission architecture is constructed.

MORT is built in a modular fashion to support its general nature. The trajectory modules for one mission (one-way CP) are likely not applicable to other missions (round trip SEP). Similarly, the spacecraft modules for a SmallSat are quite different from a crewed mission. Different missions also have different relevant constraints and mission profiles, including staging.

Given the general scope of MORT and its intended use in early mission formulation, it is critical that the time to set up and run the tool are reasonably low, otherwise it will be unwieldy to use in the dynamic concept environment of formulation. However, it must also be able to reach sufficient fidelity to give confidence in a concept. This means that MORT must be able to span the low-to-medium fidelity space.

MORT is often used in real-time design sessions as well as in extremely large tradespace explorations. This means that runtime is also a significant concern, but cannot come at the detriment of usability or development time. For the needed balance between runtime and development ease, MORT is implemented entirely in MATLAB. Runtime for a single, fully converged spacecraft and trajectory combination is typically between 0.1 to 1.0 seconds on a single core. This enables small trade space explorations to be done in real-time collaborative environments or for millions of cases to be run over a weekend on a single workstation.

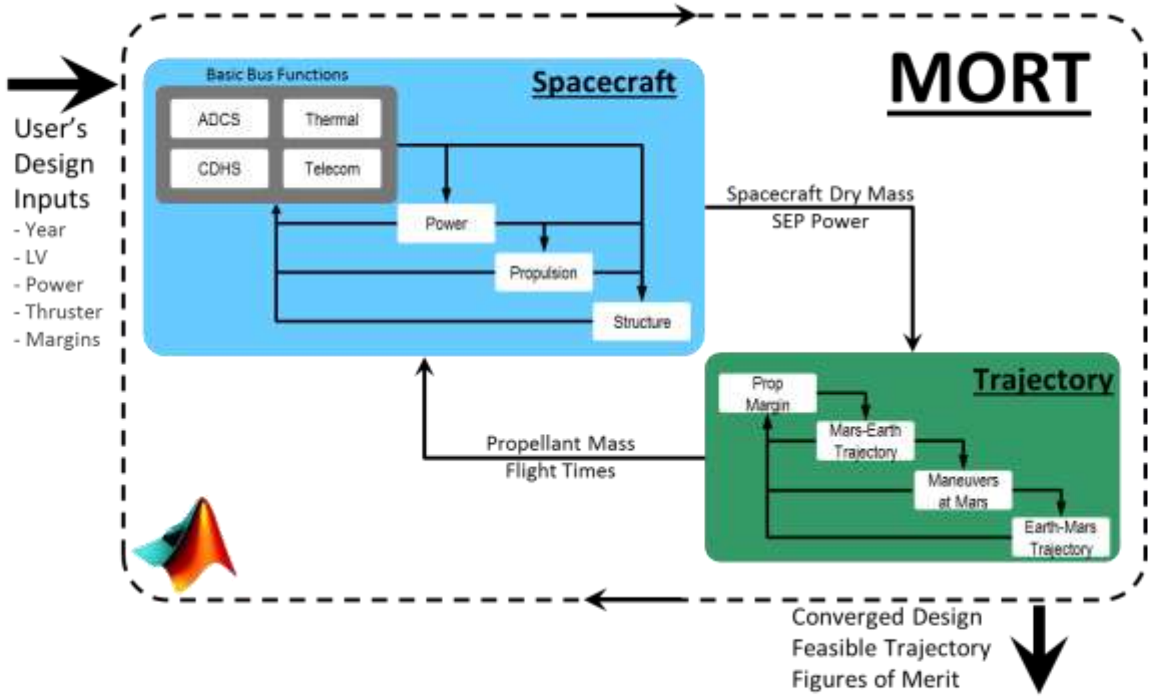

Figure 1. Simplified Block Diagram of MORT Software, Configured for Round-Trip Mars Mission

Spacecraft Module

The essence of the spacecraft module is that it is a parametric model of all of the spacecraft subsystems with respect to all variables that affect the trajectory in any meaningful way. This includes the primary sizing relationships one would expect, such as increasing solar array mass based on collecting area and increasing tank mass based on propellant, secondary relationships like longer eclipses resulting in larger batteries, and even tertiary relationships such as increased battery size resulting in higher power spent on thermal control, ultimately leading to less power available for EP and thus lower thrust.

The subsystem modules themselves can be of varying fidelity. Taking the thermal subsystem as an example, a first-order model might be a simple scaling relationship based on historical missions, such as thermal mass = 3% of dry mass. Early in concept formulation, this very general approach is likely sufficient, as more may not be known about the thermal control approach at all. More detailed parametric models might include computing radiator size and mass based on needed dissipations (which are themselves calculated based on the propulsion and spacecraft dissipation needs) and a knowledge of likely radiator technology. Based on the mission need, some subsystem models may need more detailed models than others. For instance, on SEP missions, the scaling of the solar array mass and power system is very important, while on a CP mission it is a much smaller contributor.

For some subsystems, it makes more sense to use a pick-list of component choices rather than parametric scaling relationships. A good example of this is EP engines. There are only a few of these products in the world, and they do not fall along any kind of sizing continuum. They each have a unique mass, thrust, Isp, power, and throughput limit. It may also be the case that this hardware selection is an input (due to a particular industrial preference for a given technology) instead of an output.

A primary use-case of MORT is to quantitatively study the impacts of spacecraft choices on the overall mission feasibility. Therefore, being able to quickly examine options with different subsystem choices (e.g. different EP engines, or a custom-built structure instead of a heritage one) is essential. Thus, not only is the spacecraft model modular and parametric, but it also needs to be able to switch between subsystem models with minimal user overhead in a batchable way (i.e. without needing manual reconfiguration).

At the final iteration of MORT, a full Mass Equipment List (MEL) is produced, listing all of the hardware in the converged spacecraft (at whatever level of detail modeled) and a suite of common figures of merit (such as battery capacity, solar array area, thermal loads, etc). Additional mission-specific figures of merit can be defined and reported as well.

Trajectory Module

The purpose of the trajectory module is to find an optimal trajectory meeting the user-input constraints and objectives given the mass returned by the spacecraft module. Like the spacecraft module, the trajectory module can have a variable level of detail. An extremely simplistic model would be to have a parametric model of ΔV as a function of average acceleration or some other very simple curve-fit style relationship based on previous work. However, this approach generally does not work well enough for medium fidelity concepts involving heliocentric SEP trajectories, as the nuance and complexity of the interaction between power, mass, solar distance, and planetary position has not yet been well captured by simple relationships at this time.

A higher fidelity approach is to leverage existing tools for trajectory optimization. For MORT, the Mission Analysis Low Thrust Optimizer (MALTO) tool[5] (a medium fidelity tool developed at JPL) is used for the heliocentric EP trajectories and JPL-internal tools are used for planetocentric maneuvers. A simple approach would be to plug MALTO in directly as the trajectory module and recompute the trajectory at each iteration. However, there are four main issues with this approach:

- **Runtime** – each MALTO run takes on the order of one to ten seconds, and each run takes many iterations, meaning tens of seconds in just the trajectory computation per run. This is much slower than desired (0.1 to 1.0 second for all iterations)
- **Initialization issues** – a downside of many iterative methods is that they require an initial guess. Using MALTO directly would require that the initial guess and all subsequent iterations are feasible, or at least return results that push convergence in the right direction. In practice, coming up with good initial guesses is difficult, time-consuming, and prone to human error.
- **Failure to converge or convergence to suboptimal points** – as with most processes of numerical optimization, the results produced by MALTO are not perfect. Sometimes cases fail to converge, or converge to local minima that are out of family with surrounding points. In nearly all cases, these are not true physics issues, but rather numerical issues associated with the optimization process as applied at large scale. With human tweaking of the seed case or optimization parameters, the correct solution could be retrieved, but it is time-intensive to do so and generally not possible on an iteration-by-iteration basis. If these non-optimal cases were used directly, it would likely lead to incorrect conclusions. If the cases fail to converge, the process would crash entirely and no results would be returned.

To address these problems, a database-driven interpolation method (using aforementioned MALTO and JPL tools) is usually used instead. The databases generated are a parametric sweep across the parameters to be studied over the ranges of interest. For SEP mission concepts, these typically include the launch vehicle, launch and arrival dates, EP engine types and quantities, and power levels. Depending on the mission, other parameters may also be varied. These databases are also known as “bacon plots”. Databasing is done for each relevant leg of the trajectory (e.g. Earth-Mars, Mars-Earth). The process for producing these databases is described in detail in the companion papers[6,7]. For large tradespaces, the databases can be very large. To date, over 1 billion trajectories have been databased for MORT.

However, the database creation is only part of the story. As mentioned previously, the results produced by MALTO (especially in such huge sweeps of parameters) are not perfect and sometimes include failures to converge or suboptimal convergence. In addition to these numerical issues, there are physics-based issues associated with true infeasibility (e.g. not enough acceleration to complete required ΔV in the needed time). All of these cases are eliminated before proceeding.

To correct these issues, an iterative interpolation process known as “combing” [*] is sometimes used to fill in missing datapoints and fix any obvious outliers that are out of family with nearby points. Experience has shown that SEP trajectories are suitable for interpolation because there are rarely sharp changes in behavior over small changes in the input parameters. An example of a combed dataset is shown in Figure 2. To prevent combing from inventing infeasible data, only points which can be heuristically known to be possible (e.g. with power above a known feasible power level) are filled with combing. Test cases have shown that combing generally works well.

[*] Plots of performance versus increasing power should be monotonically increasing (i.e. iso-power lines should not cross). Initial output from MALTO often does not follow this heuristic due to issues outlined and iso-power lines appear “tangled”, hence the term “combing” to smooth out and detangle the lines.

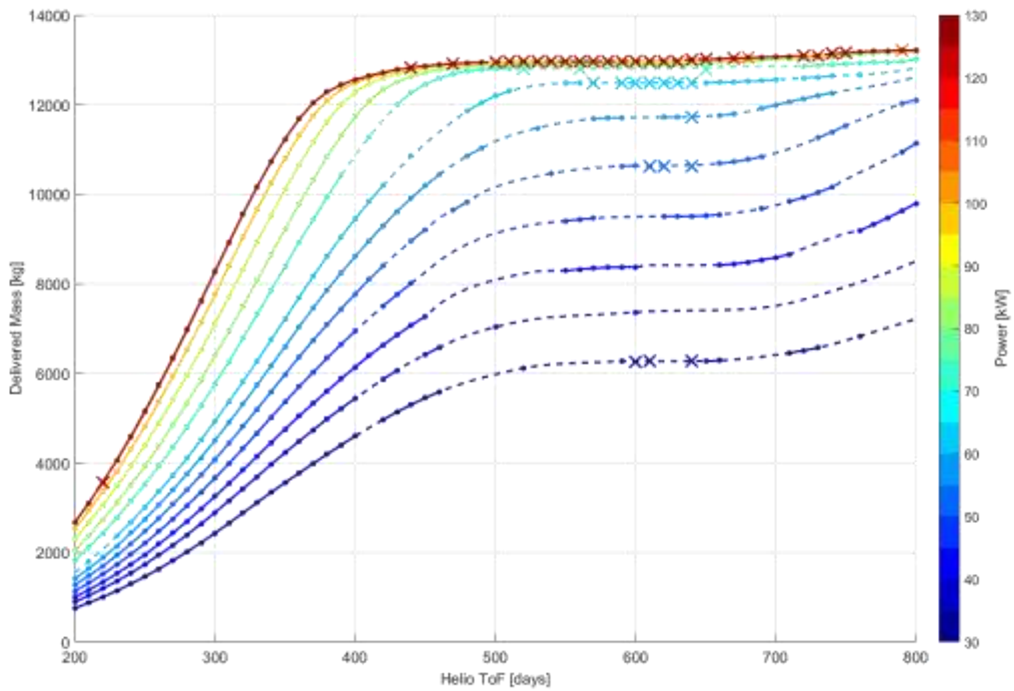

Figure 2. Example combing algorithm result. Dots are converged datapoints, solid lines connect converged datapoints. Crosses represent converged datapoints which were rejected for being out-of-family low. Dashed lines were interpolated to fill in the missing or rejected original data.

Even with reasonably large steps in the relevant dimensions, the databases can often be prohibitively large such that creating them is computationally bottlenecking for studies. To reduce the databasing compute time without being forced to consider only a small number of discrete values in each dimension, interpolation is used instead of a simple lookup. In this case, a kriging approach was used, based on the DACE[8] implementation. Use of kriging interpolants is a common technique in computer experiments, used to create surrogate models based on limited output. In this case, this technique is used as a surrogate for data between grid points or that is missing due to numerical issues as noted above. This is especially important when applying constraints or filters to data and it can become somewhat sparse or irregular, and accurate prediction is important. Kriging was found to be superior to other techniques in accuracy of interpolation, ease of utilization (due to being able to handle few datapoints with irregular spacing) and in validity of extrapolations (more on this in subsequent paragraphs). Using interpolation in this way allows a dramatic reduction of database size and also allows the modeling of arbitrary values within the range of interest rather than a pre-fixed set. Verification activities have shown that interpolations of this type are very accurate, typically within 1% of the value which would be obtained with a dedicated MALTO run.

A key benefit of this database-driven interpolation approach is that it allows the trajectory to be effectively broken into leg-databases and mass-matched in MORT, rather than having to database all possibilities of the entire mission. For instance, consider a mission of two major propulsive legs (such as a round-trip crewed mission), where the first leg has N available date options, the second leg has N available date options, the spacecraft has P available propulsion options, L available launch vehicles, M available mass options at the end of the second leg, and S staging options between the two legs. This would result in N^2PLMS runs to fully enumerate the space. Typically N >> L,P,M,S. The databasing approach used in MORT instead databases each leg separately and uses interpolation to mass-match the legs. This means the total database size is NLP+NPM. Note that the staging variable is entirely eliminated, as this can be added between the legs without any databasing at all. Doing leg-based databases instead of end-to-end trajectory databases ultimately reduces the number of database points by a factor of ~10^6, which is absolutely critical for the feasibility of this type of approach. For missions with more than two legs to be databased, the benefit is even higher. From a day-to-day use perspective, this approach enables mixing and matching of databases to create combinations not originally conceived ahead of time.

This interpolation method also addresses the initialization issue mentioned previously. MORT deals with this problem by ensuring that the interpolants are capable of extrapolation with the correct trend (e.g. increasing propellant load for increasing mass). In this way, even if the initial guess is not good, it will yield an extrapolated result and, over the subsequent iterations, will either enter into the interpolation domain (which are actually feasible), or converge outside of it, in which case the result is deemed infeasible. In this way, MORT can be guaranteed to always output a quantitative result which is known to be either feasible or infeasible. Verification activities have shown that this approach is extremely robust to very poor initial guesses, with essentially no relevant impact on the converged results. For this reason, there is no longer a need for a user-input initial guess at all – a simple algorithm based on other inputs (payload, power level, etc) is used to seed a low-quality initial guess and iterations proceed from there to convergence. Removing the need for a user-input guess is a great boon to usability, as for some numerical methods, coming up with the initial guess is the most time-consuming part of the process.

The trajectory module returns a number of important figures of merit, the most important of which are the mission timeline (dates of all relevant events and maneuvers) and the maneuver list (which includes a time history of mass, ΔV, and propellant from launch to end of mission). Each mission also has an optimized launch vehicle target (C3, declination, azimuth) based on the optimal trajectory. For EP missions, relevant intermediate EP figures of merit like ΔV per day and instantaneous acceleration are also reported. As in the spacecraft module, mission-specific figures of merit can be defined and reported. Because the trajectory model is interpolated, a true MALTO output is not generated, but the MORT outputs can be used to quickly re-create the MALTO case. This has been done many times for a variety of concepts and is generally a smooth process because the MORT interpolation process works well.

APPLICATION EXAMPLES

To demonstrate the utility that MORT provides, three application examples are shown. In order to provide a focused paper, only select analysis results are shown. The goal is not to dwell on the specifics of the mission concepts or the quantitative input parameters, but rather the useful conclusions that MORT is able to offer that would be otherwise unavailable. Additional work performed using MORT can be found in the companion paper[9] and also these other references[1,3,4].

Remote Science Orbiter Example

The first application considered is for a low Mars orbit remote science orbiter concept, somewhat similar to most previous NASA orbiters such as Mars Odyssey (ODY) and Mars Reconnaissance Orbiter (MRO). Relevant input assumptions were that the mission would carry 200kg of scientific payloads and should be compatible with a reusable launch vehicle (for which Falcon 9 performance models were used as a proxy). The science goals for this mission involved taking measurements from multiple altitudes and multiple inclinations, which ruled out CP + aerobraking approaches like those used on MRO and ODY as propellant-infeasible. Therefore, SEP approaches were the only ones considered. The primary maneuvers are the heliocentric transfer from Earth to Mars with a C3≈0 approach, followed by a nearly circular spiral down to Low Mars Orbit. Options were studied for launch in 2022 and 2024, but for EP missions the sensitivity to launch opportunity has been found to be low. Total ΔV is in the 5-7 km/s range.

MORT was used to assess the feasibility and optimality of the various thruster options for this mission application. The thrusters considered are listed in Table 1. Note that the SPT-140 thruster produced by Fakel was also considered, but performance was similar enough to XR-5 that additional cases were not run. Results from XR-5 cases can be considered representative of SPT-140 at this level of analysis. Table values are given at maximum power for context; performance curves as a function of available power are used in the analysis. The performance differences in the multiple-mode thrusters (XR-5, NEXT) are more pronounced at lower power. Lastly, it is worth noting that these values are used for analysis purposes and may differ somewhat from the exact capabilities of the systems.

Table 1. List of EP engines and models used for MORT analysis of Remote Science Orbiter Example

Thruster Name	Developer	Mode	Code in Plots	Power to PPU [kW]	Thrust [mN]	Isp [s]	Thrust / Power [mN/kW]	String Mass [kg]
XR-5 (BPT4000)	Aerojet	Higher Thrust	BPT	4.8	281	1865	58	34
		Higher Isp	BPTi	4.8	252	2010	52	34
XIPS	L-3	-	XIPS	5.0	173	3507	34	50
NEXT-C	NASA / Aerojet	Higher Thrust	NEXT	6.9	227	3999	33	84
		Higher Isp	NEXTi	6.9	219	4077	32	84
HERMeS	NASA / Aerojet	-	ARMv3	13.5	612	2805	45	114

The MORT output addressing launch feasibility is shown in Figure 3. The chosen figure of merit is launch vehicle margin (LVM). This represents the mass utilization of the launch vehicle at the optimized C3 value. Negative values indicate infeasibility. Figure 4 shows the associated time of flight from Earth launch to arrival at LMO. Each line constitutes a sweep over solar array sizes for a given engine configuration. Each point along each line represents a fully converged spacecraft with a fully constructed equipment list and a matching trajectory which is feasible in mass and time. The spacecraft itself is sized according to a parametric model and is re-sized to accommodate the needed propellant loads, solar arrays, thrusters, etc. This represents around 700 spacecraft designs considered, which is far more than could be considered in equivalent detail without a tool like MORT. There is some "noise" in the results due to various numerical imperfections in the source MALTO data, various effects in the spacecraft model, and numerical errors in the MORT iterative process. However, the trends are clearly visible and conclusions can be drawn. Legend nomenclature is [EngineCode]x[ActiveThrusterQty] (+[SpareThrusterQty]). All mission concepts carry one "cold" spare for redundancy "(+1)". Because EP engines also have a maximum usable throughput (often measured as total impulse, in mN), some concepts carry additional spares to compensate.

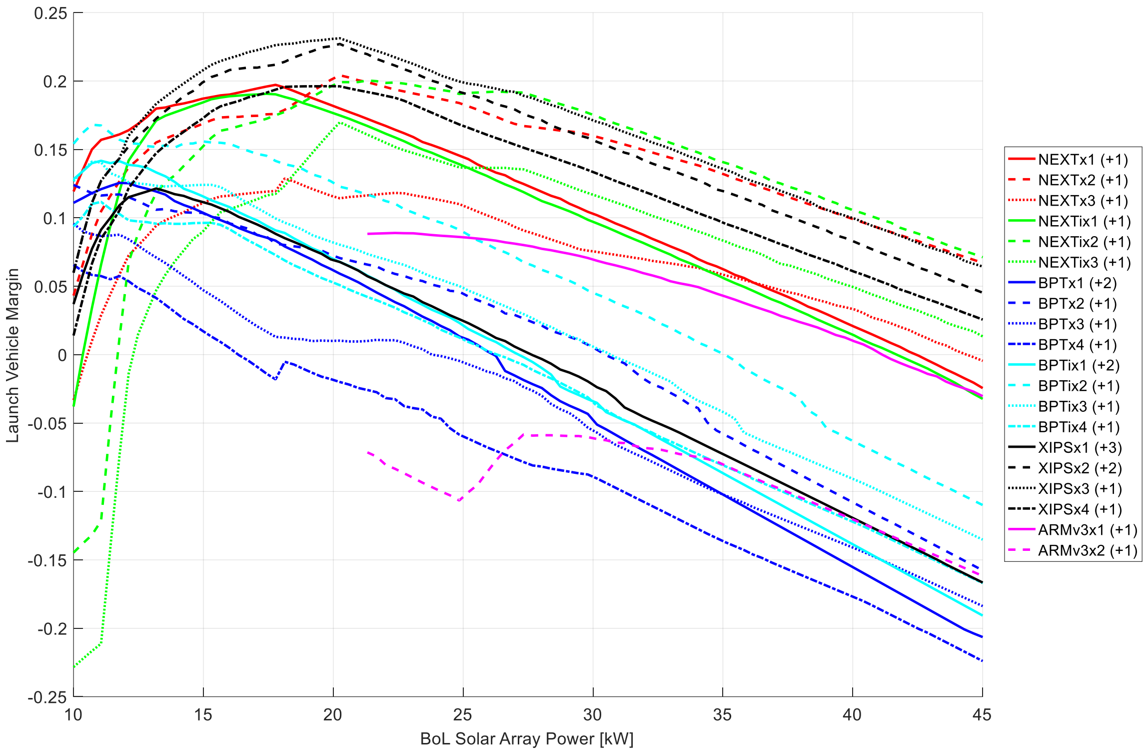

Figure 3. Launch vehicle margin as a function of solar array power and EP engine choice for Remote Science Orbiter example case

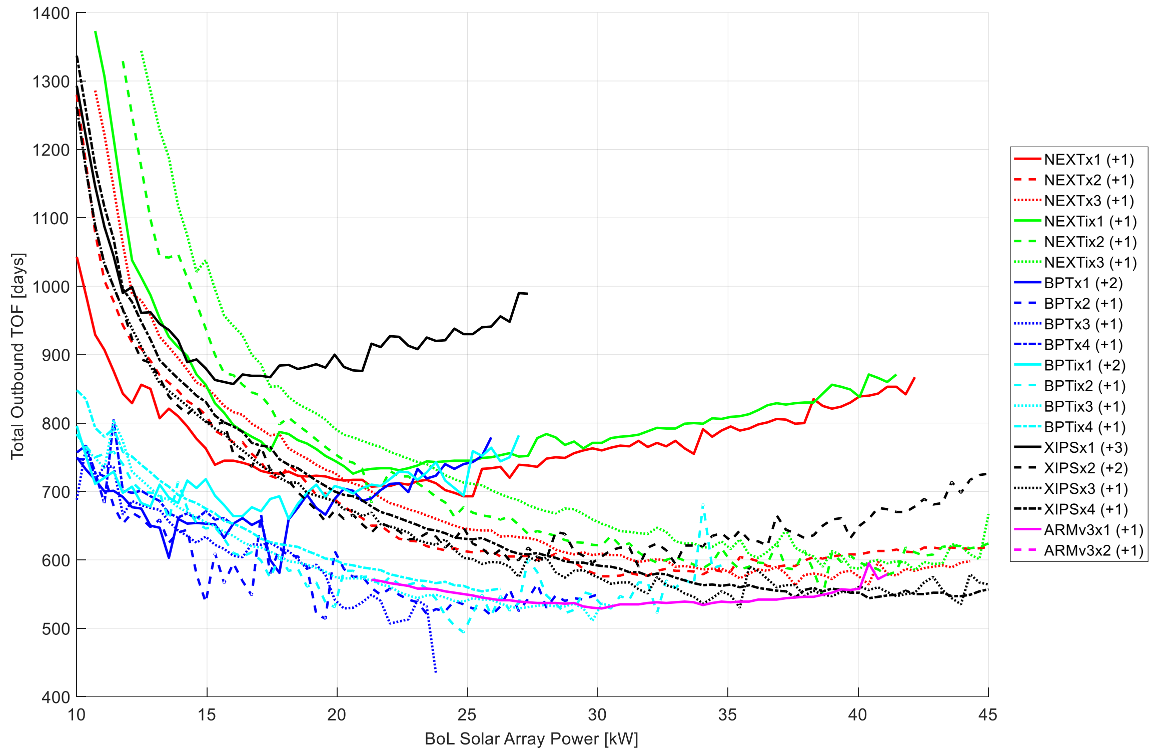

Figure 4. Total Time of Flight (launch to arrival in Low Mars Orbit) as a function of solar array power and EP engine choice for Remote Science Orbiter example case. Only solutions with positive launch vehicle margin are shown, and all solutions are matched to the LVM-optimized results shown in Figure 3.

From these plots, several useful conclusions can be drawn:

- The mission is shown to be feasible within the launch constraints using any of the studied engine technologies as long as the solar array size and quantity of active engines can be chosen appropriately. Therefore, other factors such as flight time, cost, etc, can be the basis for this choice.
- For each engine choice, there is a power level which is optimal from a launch vehicle margin (LVM) perspective. This is the breakpoint at which adding the mass of additional solar arrays is exactly balanced with the additional propulsive benefit that power adds. Additional power beyond this point may still be useful in other figures of merit (for example, by reducing time of flight). To the left of this optimal power point, the launch vehicle margin typically falls off rapidly while to the right of this optimal power point, the launch vehicle margin typically falls off more slowly.
    - The XR-5 options have the lowest LVM-optimal power point, around 10 kW. Lower power values were excluded due to science payload needs but the true optimal could be even lower.
    - The NEXT and XIPS options have LVM-optimal power points around 20 kW but it is a shallow maximum and powers in the range of 15 - 25 kW generally work well.

- The HERMeS option needs more than 20 kW to function at the minimum operating point for a single thruster (roughly 50% of max power) at Mars range. Additional power beyond this is not very helpful.

- While each engine has some unique attributes, the major differences between options is based on the EP engine type: Hall Effect (XR-5, HERMeS) or Gridded Ion (NEXT, XIPS). This is a very natural conclusion given that these two thruster types differ by a factor of approximately two in thrust and Isp.
- The Hall Engines typically have lower launch vehicle margin (though still positive) and lower times of flight as compared with the Gridded Ion options at a given power level.
- For the XIPS configuration, it would be very beneficial to have at least two active engines. For the HERMeS configuration, no more than 1 active engine is desirable. For the other configurations, there is a modest sensitivity of LVM and TOF to number of engines but most combinations could work well.
- For engines with multiple modes, there is a modest sensitivity to higher thrust vs higher Isp mode, with the former having lower flight time and higher LVM.
- Flight times from Earth to Mars LMO can generally be expected to be between 500 and 900 days, with the LVM-optimal configurations generally being around 700 days. Additional power beyond the LVM-optimal can decrease flight time.
- (From additional plots omitted from this report for brevity) Expected values for other parameters are in the following ranges:
    - Xenon Mass: 700-1000kg (Hall Effect) or 300-600kg (Gridded Ion)
    - Spacecraft Dry Mass: 1500 – 2000 kg
    - C3: 4 – 8 km^2/s^2 (Hall Effect) or 6 – 12 km^2/s^2 (Gridded Ion)
    - Nominal launch date within a given opportunity is not strongly a function of engine or power, with around 1 month of variation across all options studied. Therefore, the time of flight results in Figure 4 are essentially the same as what an arrival date plot would show.

These general trends, as well as the more detailed quantitative information which can be retrieved for individual runs, has been found to be immensely useful in the often dynamic environment of concept formulation. Of course, there are many more aspects other than trajectory feasibility which ultimately affect the choice of EP technology, with cost and heritage being two of the more important ones. MORT enables these decisions to be made with much better quantitative assessment of the potential impacts than could have been done in the past.

Mars Sample Return Multifunction Orbiter Example

The second example mission considered is a multifunction orbiter with both in-situ remote science objectives as well as Mars Sample Return (MSR) objectives, similar to the Next Mars Orbiter (NeMO) concept based on the Next Orbiter Science Analysis Group (NEX-SAG) report [10]. Additional background on MSR mission concepts can be found in the companion paper[11] and in other older references[12,13] and will not be addressed further here. From a trajectory perspective, the most important difference is that this concept is a round-trip (return to Earth) option, which more than doubles the required delta V (13 – 15 km/s). This also means that the MORT results must consider the competing impacts of mass, power, etc on both legs of the transfer, and that the Mars-Earth leg happens entirely under the propulsive capability of the orbiter (no launch vehicle boost). Total mission time was expected to be around 9 years (~3300 days), including all transfers, science time at Mars, and MSR rendezvous operations.

To reduce the return mass, it is assumed that 66% of the payload equipment is jettisoned before the return leg. For the purpose of this analysis, Falcon Heavy (with re-use) capability was assumed, but results would be similar or better for Falcon 9 (expendable), Atlas V 551, Vulcan, New Glenn, and others of similar performance. The payload in this case would include not only remote sensing equipment, but also the sample capture payload and perhaps additional daughter-craft with science or communications objectives. A payload mass around 200kg would be considered the bare minimum (e.g. accomplishing minimum science and MSR objectives, with focus on mass reduction) with more being better. The time spent in LMO before departing for Earth return is an important figure of merit, as this is the time available to use the science payload.

For this example, two optimization approaches are considered. The first is similar to the previous section, where a payload mass of 200kg is assumed. This is termed the "minimum payload" approach. An additional analysis ("maximum payload") is also shown, where instead of fixing the payload mass and measuring launch vehicle margin, the payload mass was optimized within the launch vehicle constraint.

The MORT results are shown in subsequent plots. Because the payload mass is an input to MORT, solving the maximum payload requires wrapping MORT in a one-dimensional optimizer solving for payload mass. This represents around 700 spacecraft designs output, with around 20,000 or so considered during the payload optimization process. This is far more than could be considered in equivalent detail without a tool like MORT. Again, there is some "noise" in the results but the trends are evident.

Figure 5, Figure 6, and Figure 8 are associated with the max payload case, while Figure 7 is associated with the minimum payload case. Thruster types and legends follow the same form as previous figures. Figure 5 shows the sensitivity of useful payload mass to solar array sizing and engine type. All cases shown have positive launch vehicle margin. Figure 6 and Figure 7 show the total time of flight (includes Earth-Mars and Mars-Earth heliocentric legs and the altitude raising and lowering spirals). Because the total mission time does not vary much, these plots can also be used to find the stay time at Mars by subtracting the values from the total mission time of around 3300 days.

Figure 8 shows a more detailed mass breakdown for the case of two higher-Isp-mode active NEXT thrusters. The mass is categorized by subsystem and plotted as a function of solar array size. The most notable trends as the solar array size increases are the increase in total mass (due to trajectory improvements), propellant use (as total mass grows, and also as the propulsive burden shifts more toward the spacecraft from the launch vehicle) and power subsystem (mass increases with solar array size).

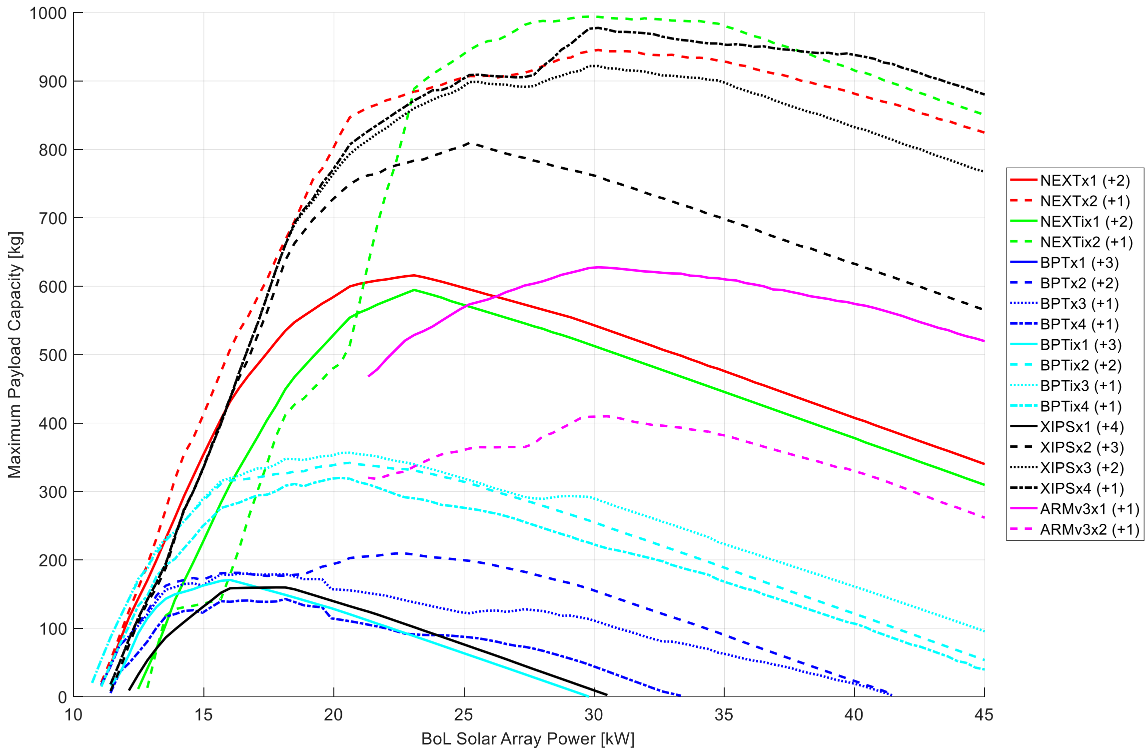

Figure 5. Example MORT result showing sensitivity of useful payload mass to solar array sizing and EP thruster choice for the MSR SEP example

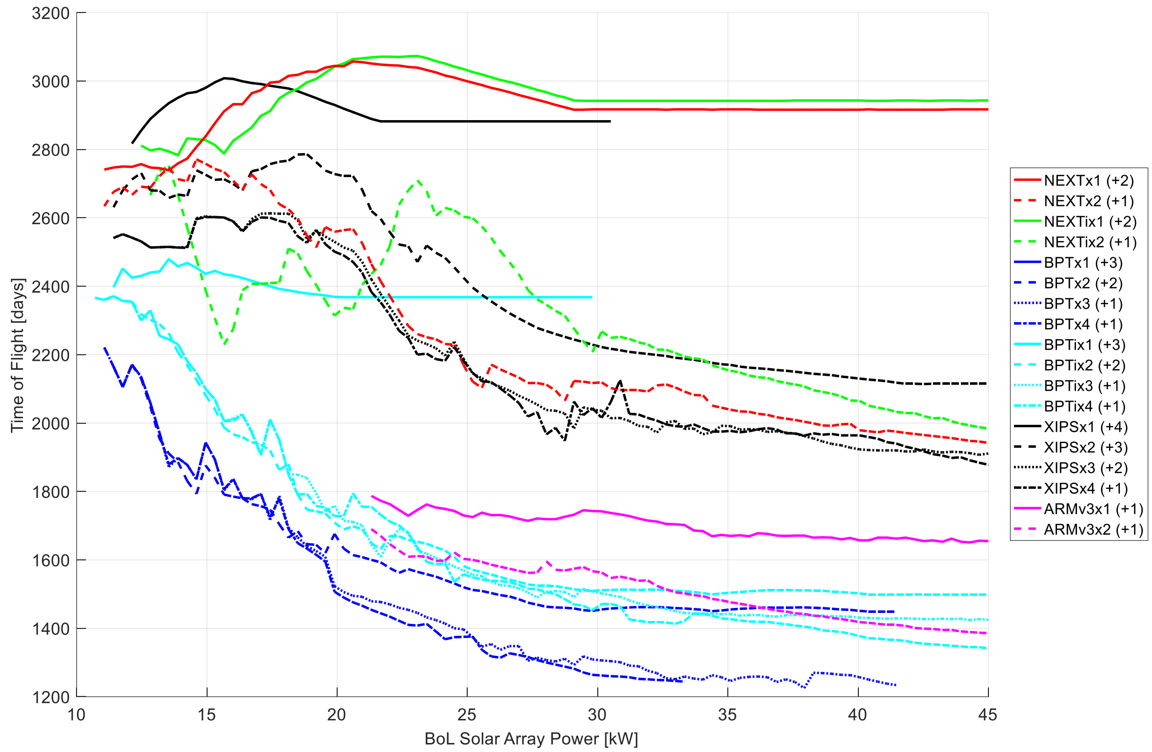

Figure 6. Total time of flight (all legs) as a function of solar array power and EP engine choice for the MSR SEP Example case, with optimized payloads. All solutions are matched to the LVM-optimized results shown in Figure 5. Launch and return dates are not very sensitive to power or engine type, so the time in LMO can be computed from this plot by subtracting it from the total mission time.

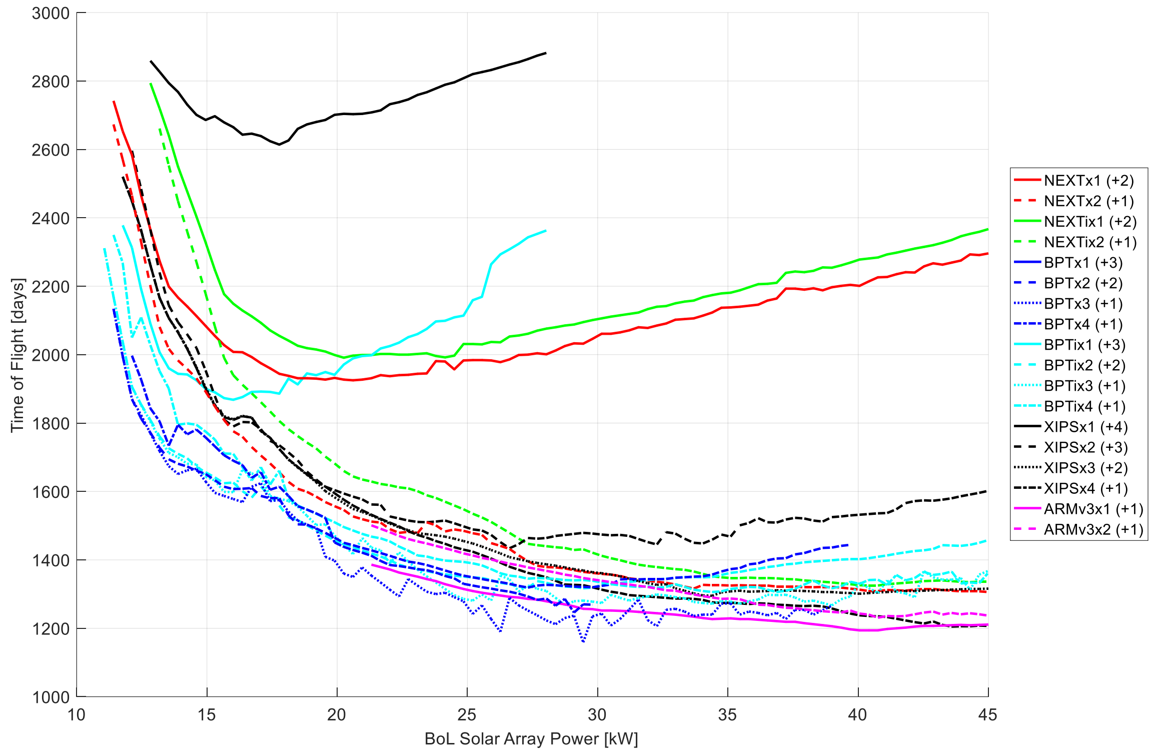

Figure 7. Total time of flight (all legs) as a function of solar array power and EP engine choice for the MSR SEP Example case, with fixed payload mass of around 200kg and positive launch vehicle margin. These results are NOT matched with Figure 5. Launch and return dates are not very sensitive to power or engine type, so the time in LMO can be computed from this plot by subtracting it from the total mission time.

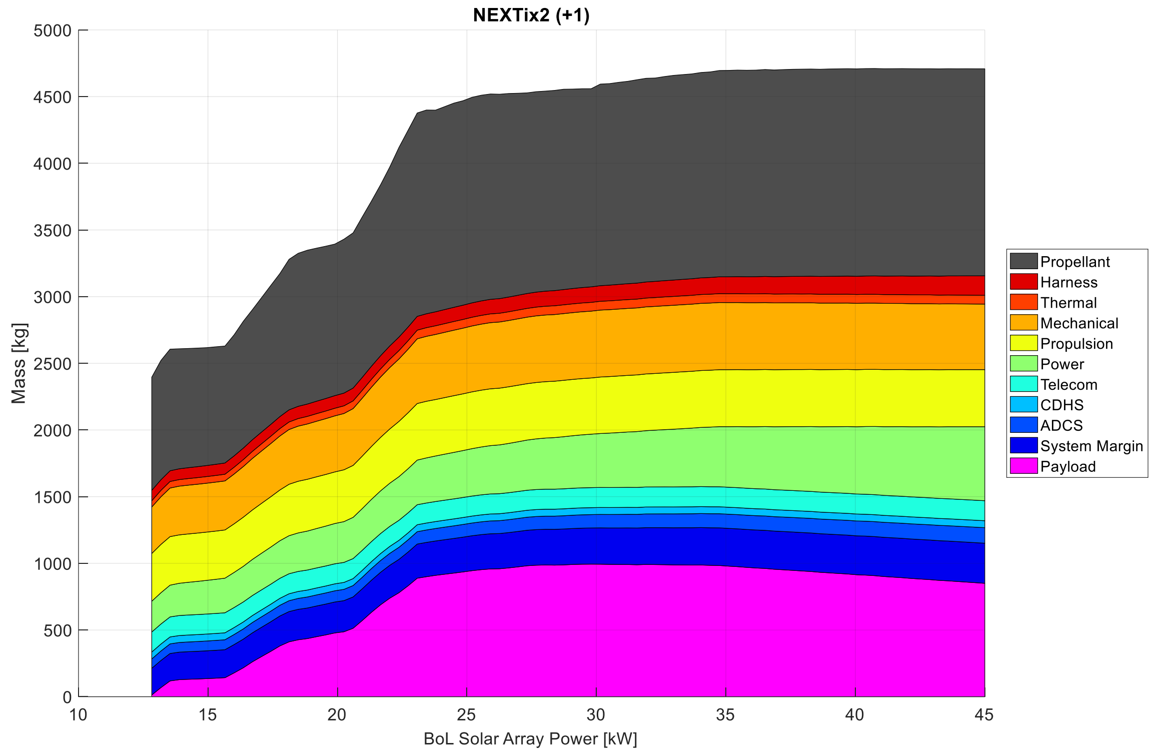

Figure 8. Mass breakdown by subsystem for NEXTix2(+1) case from Figure 5. The magenta contour is the same as the dashed green line.

From these plots, many useful conclusions can be drawn:

- For each engine type, there are solutions yielding maximum payloads above the minimum around 200kg, as long as the solar array size, engine mode, and number of active engines can be chosen appropriately.
- As with the Remote Science Orbiter example, the same behavior of an optimal power level is once more present in these results, with sharp drops as power decreases and shallow drops as power increases
  - For the XR-5 option, this optimal solar array size is around 20 kW
  - For the other thruster options, the optimal power is around 30 kW
- The two gridded ion options (NEXT, XIPS) show very similar performance, but the two hall effect options (XR-5, HERMeS) are noticeably different. This is likely because the XR-5 option carries additional engines to meet throughput needs, reducing its mass efficiency, while the single HERMeS engine is capable of doing the entire mission on a single engine.
- The higher thrust mode XR-5 option is the only thruster/mode combination which is marginal on the minimum payload mass limit. The higher Isp mode of XR-5 can achieve over 300kg payload across a reasonable range of solar array sizes.
- There would be a very significant benefit in having at least two active NEXT, XIPS, or XR-5 engines, either to increase payload mass or decrease time spent in transit.

Additional incremental performance is available for XIPS and XR-5 in a third active engine. If using the HERMeS engine technology, one active engine is preferred.

- When maximizing payload, the gridded ion options are capable of much higher maximum payloads than the hall effect options, especially as compared to XR-5. However, this comes at the price of a significantly longer flight time, reducing the amount of time that payload can be used in LMO.
- When holding the payload fixed at its minimum value, the gridded ion options are slightly slower than the hall effect options at a given power level (up to 200 days). However, this could be offset by increasing the power by around 5 kW if desired. Ultimately, there is likely little difference between these options in the minimum payload case from a trajectory perspective.
- Based on the above two conclusions, from a transit perspective, gridded ion technology is preferred in all cases.
- In the minimum payload case, the total transit time if operating around the optimal power levels is around 1500 days. This means roughly 5 years would be available in LMO for science operations. As the payload increases, the time at Mars decreases. For the gridded ion options with the maximum payload potential (nearly 1000kg), time at Mars could be reduced to 3 years or even lower depending on chosen power level.
- (From additional plots omitted from this report for brevity) Expected values for other parameters are in the following ranges. Ranges are large due to the large variation in power and payload masses considered:
    - Xenon Mass: 1700 – 2400 kg (Hall Effect) or 800 – 1300 kg (Gridded Ion)
    - Spacecraft Dry Mass: 2000kg – 3000 kg
    - C3: 2 – 16 km^2/s^2
    - Nominal launch date and Earth return date within a given opportunity is not strongly a function of engine or power, with around 1 month of variation across all options studied.

As in the previous example, the MORT analysis offers conclusions which would be otherwise difficult to obtain without the ability to convolve trajectory and system design. This is valuable information in concept formulation. In particular, a powerful conclusion is that a spacecraft with 2-3 active NEXT or XIPS engines with a power level around 30 kW is optimal for maximizing payload delivery and also offers a low-time of flight for lighter payloads. This optimal point would be difficult to determine intuitively due to all the competing factors.

While it is not surprising that the higher Isp choice can offer the most payload, it is counter-intuitive that it can also have a similar stay time at Mars for lower payloads. One would expect the hall effect options to be able to achieve a significantly lower time of flight due to increased thrust. However, because of the mass saved by the higher Isp, the gridded ion options can afford a higher power level, which ultimately leads to a similar flight time. Further, the hall effect options' accelerations are decreased by the propellant mass needed to perform the large delta V, especially on the Earth-Mars leg.

CONCLUSIONS

The capabilities of MORT have been instrumental in both framing the mission trade space and performing specific impact sensitivity studies, allowing the discovery of non-intuitive design solutions and quickly find optima that might never be otherwise considered in a more traditional sequence-of-point-designs approach. A selection of examples was shown to illustrate the utility of MORT in deriving system-level conclusions based on the interplay between spacecraft and trajectory and how these can be optimized together.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The information presented about potential Mars Sample Return is pre-decisional and is provided for planning and discussion purposes only.

REFERENCES

[1] R.C. Woolley, J.D. Baker, D.F. Landau, and A.K. Nicholas, "Cargo Logistics for a Notional Mars Base Using Solar Electric Propulsion," *Acta Astronautica*. In Press, accepted Aug. 13, 2018. https://doi.org/10.1016/j.actaastro.2018.08.026

[2] S. Matousek, A.B. Sergeyevsky, "To Mars and Back: 2002-2020 Ballistic Trajectory Data for the Mission Architect", *AIAA/AAS Astrodynamics Specialist Conference*, A98-37376, AAS, 1998. Web

[3] Woolley, R. C., Nicholas, A. K. "SEP Mission Design for Mars Orbiters," *AAS/AIAA Astro-dynamics Specialist Conference*, AAS15-632, AAS, 2015. Web.

[4] Snyder, J. S. et. al. "Additional mission applications for NASA's 13.3-kW Ion propulsion system," *Aerospace Conference, 2016 IEEE*, IEEE, 2016. Web.

[5] J. A. Sims, P. A. Finlayson, E. A. Rinderle, M.A. Vavrina, and T. D. Kowalkowski, "Implementation of a Low-Thrust Trajectory Optimization Algorithm for Preliminary Design," AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Keystone, Colorado, Aug. 2006.

[6] F. Laipert, A.K. Nicholas, Z.P. Olikara, R.C. Woolley, and R.E. Lock, "Hybrid Chemical-Electric Trajectories for a Mars Sample Return Orbiter," AAS Astrodynamics Specialist Conference. Maui, HI, Jan. 2019.

[7] R.C. Woolley, F. Laipert, A.K. Nicholas, and Z.P. Olikara, "Low-Thrust Trajectory Bacon Plots for Mars Mission Design," AAS Astrodynamics Specialist Conference. Maui, HI, Jan. 2019.

[8] Nielsen, H.B., Lophaven, S.N., Søndergaard, J. "*DACE – A Matlab Kriging Toolbox*," Informatics and Mathematical Modelling, Technical University of Denmark, DTU. 2002. Web.

[9] A.K. Nicholas, A. Didion, F. Laipert, Z. Olikara, R.C. Woolley, R.E. Lock, "Mission Analysis for a Potential Mars Sample Return Campaign in the 2020's," AAS Astrodynamics Specialist Conference. Maui, HI, Jan. 2019.

[10] MEPAG NEX-SAG Report, Report from the Next Orbiter Science Analysis Group (NEX-SAG), Chaired by B. Campbell and R. Zurek, Mars Exploration Program Analysis Group (MEPAG), 2015. Web.

[11] R.E. Lock, A.K. Nicholas, S. Vijendran, R.C. Woolley, A. Didion, F. Laipert, Z. Olikara, "Potential Campaign Architectures and Mission Design Challenges for Near Term International Mars Sample Return Mission Concepts," AAS Astrodynamics Specialist Conference. Maui, HI, Jan. 2019.

[12] Lock, R. E. et. al. "Mars Sample Return Orbiter concepts using Solar Electric Propulsion for the post-Mars2020 decade," IEEE Aerospace Conference, 2014. Web.

[13] R. Mattingly, L. May, "Mars Sample Return as a Campaign", IEEE Aerospace Conference, 2011. Web.